

**SYSTEMS, METHODS, AND SOFTWARE FOR CLASSIFYING TEXT
FROM JUDICIAL OPINIONS AND OTHER DOCUMENTS**

Related Application

5 This application is a continuation of U.S. Provisional Application
60/336,862, which was filed on November 2, 2001 and which is incorporated
herein by reference.

Copyright Notice and Permission

10 A portion of this patent document contains material subject to copyright
protection. The copyright owner has no objection to the facsimile reproduction
by anyone of the patent document or the patent disclosure, as it appears in the
Patent and Trademark Office patent files or records, but otherwise reserves all
copyrights whatsoever. The following notice applies to this document:
Copyright © 2001, West Group.

15

Technical Field

The present invention concerns systems, methods, and software for
classifying text and documents, such as headnotes of judicial opinions.

20

Background

25 The American legal system, as well as some other legal systems around
the world, relies heavily on written judicial opinions ---the written
pronouncements of judges--- to articulate or interpret the laws governing
resolution of disputes. Each judicial opinion is not only important to resolving a
particular legal dispute, but also to resolving similar disputes in the future.
Because of this, judges and lawyers within our legal system are continually
researching an ever-expanding body of past opinions, or case law, for the ones
most relevant to resolution of new disputes.

30 To facilitate these searches, companies, such as West Publishing
Company of St. Paul, Minnesota (doing business as West Group), not only
collect and publish the judicial opinions of courts across the United States, but
also summarize and classify the opinions based on the principles or points of law

they contain. West Group, for example, creates and classifies headnotes --short summaries of points made in judicial opinions-- using its proprietary West Key Number™ System. (West Key Number is a trademark of West Group.)

5 The West Key Number System is a hierarchical classification of over 20 million headnotes across more than 90,000 distinctive legal categories, or classes. Each class has not only a descriptive name, but also a unique alpha-numeric code, known as its Key Number classification.

10 In addition to highly-detailed classification systems, such as the West Key Number System, judges and lawyers conduct research using products, such as American Law Reports (ALR), that provide in-depth scholarly analysis of a broad spectrum of legal issues. In fact, the ALR includes about 14,000 distinct articles, known as annotations, each teaching about a separate legal issue, such as double jeopardy and free speech. Each annotations also include citations and/or headnotes identifying relevant judicial opinions to facilitate further legal
15 research.

To ensure their currency as legal-research tools, the ALR annotations are continually updated to cite recent judicial opinions (or cases). However, updating is a costly task given that courts across the country collectively issue hundreds of new opinions every day and that the conventional technique for
20 identifying which of these cases are good candidates for citation is inefficient and inaccurate.

In particular, the conventional technique entails selecting cases that have headnotes in certain classes of the West Key Number System as candidates for citations in corresponding annotations. The candidate cases are then sent to
25 professional editors for manual review and final determination of which should be cited to the corresponding annotations. Unfortunately, this simplistic mapping of classes to annotations not only sends many irrelevant cases to the editors, but also fails to send many that are relevant, both increasing the workload of the editors and limiting accuracy of the updated annotations.

30 Accordingly, there is a need for tools that facilitate classification or assignment of judicial opinions to ALR annotations and other legal research tools.

Summary of Exemplary Embodiments

To address this and other needs, the present inventors devised systems, methods, and software that facilitate classification of text or documents according to a target classification system. For instance, one exemplary system aids in classifying headnotes to the ALR annotations; another aids in classifying headnotes to sections of American Jurisprudence (another encyclopedic style legal reference); and yet another aids in classifying headnotes to the West Key Number System. However, these and other embodiments are applicable to classification of other types of documents, such as emails.

More particularly, some of the exemplary systems classify or aid manual classification of an input text by determining a set of composite scores, with each composite score corresponding to a respective target class in the target classification system. Determining each composite score entails computing and applying class-specific weights to at least two of the following types of scores:

- a first type based on similarity of the input text to text associated with a respective one of the target classes;
- a second type based on similarity of a set of non-target classes associated with the input text and a set of non-target classes associated with a respective one of the target classes;
- a third type based on probability of one of the target classes given a set of one or more non-target classes associated with the input text; and
- a fourth type based on a probability of the input text given text associated with a respective one of the target classes.

These exemplary systems then evaluate the composite scores using class-specific decision criteria, such as thresholds, to ultimately assign or recommend assignment of the input text (or a document or other data structure associated with the input text) to one or more of the target classes.

Brief Description of Drawings

- Figure 1 is a diagram of an exemplary classification system 100 embodying teachings of the invention, including a unique graphical user interface 114;
- 5 Figure 2 is a flowchart illustrating an exemplary method embodied in classification system 100 of Figure 1;
- Figure 3 is a diagram of an exemplary headnote 310 and a corresponding noun-word-pair model 320.
- Figure 4 is a facsimile of an exemplary graphical user interface 400 that
10 forms a portion of classification system 100.
- Figure 5 is a diagram of another exemplary classification system 500, which is similar to system 100 but includes additional classifiers; and
- 15 Figure 6 is a diagram of another exemplary classification system 600, which is similar to system 100 but omits some classifiers.

Detailed Description of Exemplary Embodiments

This description, which references and incorporates the above-identified Figures, describes one or more specific embodiments of one or more inventions.

20 These embodiments, offered not to limit but only to exemplify and teach the one or more inventions, are shown and described in sufficient detail to enable those skilled in the art to implement or practice the invention. Thus, where appropriate to avoid obscuring the invention, the description may omit certain information known to those of skill in the art.

25 The description includes many terms with meanings derived from their usage in the art or from their use within the context of the description. However, as a further aid, the following exemplary definitions are presented.

The term “document” refers to any addressable collection or arrangement of machine-readable data.

30 The term “database” includes any logical collection or arrangement of documents.

The term “headnote” refers to an electronic textual summary or abstract concerning a point of law within a written judicial opinion. The number of headnotes associated with a judicial opinion (or case) depends on the number of issues it addresses.

5

Exemplary System for Classifying Headnotes to American Legal Reports

Figure 1 shows a diagram of an exemplary document classification system 100 for automatically classifying or recommending classifications of electronic documents according to a document classification scheme. The exemplary embodiment classifies or recommends classification of cases, case citations, or associated headnotes, to one or more of the categories represented by 13,779 ALR annotations. (The total number of annotation is growing at a rate on the order of 20-30 annotations per month.) However, the present invention is not limited to any particular type of documents or type of classification system.

Though the exemplary embodiment is presented as an interconnected ensemble of separate components, some other embodiments implement their functionality using a greater or lesser number of components. Moreover, some embodiments intercouple one or more the components through a local- or wide-area network. (Some embodiments implement one or more portions of system 100 using one or more mainframe computers or servers.) Thus, the present invention is not limited to any particular functional partition.

System 100 includes an ALR annotation database 110, a headnotes database 120, and a classification processor 130, a preliminary classification database 140, and editorial workstations 150.

ALR annotation database 110 (more generally a database of electronic documents classified according to a target classification scheme) includes a set of 13,779 annotations, which are presented generally by annotation 112. The exemplary embodiment regards each annotation as a class or category. Each annotation, such as annotation 112, includes a set of one or more case citations, such as citations 112.1 and 112.2.

Each citation identifies or is associated with at least one judicial opinion (or generally an electronic document), such as electronic judicial opinion (or case) 115. Judicial opinion 115 includes and/or is associated with one or more

headnotes in headnote database 120, such as headnotes 122 and 124. (In the exemplary embodiment, a typical judicial opinion or case has about 6 associated headnotes, although cases having 50 or more are not rare.) A sample headnote and its assigned West Key Number class identifier are shown below.

Exemplary Headnote:

In an action brought under Administrative Procedure Act (APA), inquiry is twofold: court first examines the organic statute to determine whether Congress intended that an aggrieved party follow a particular administrative route before judicial relief would become available; if that generative statute is silent, court then asks whether an agency's regulations require recourse to a superior agency authority

Exemplary Key Number class identifier:

15AK229 - ADMINISTRATIVE LAW AND PROCEDURE - SEPARATION OF ADMINISTRATIVE AND OTHER POWERS - JUDICIAL POWERS

In database 120, each headnote is associated with one or more class identifiers, which are based, for example, on the West Key Number Classification System. (For further details on the West Key Number System, see West's Analysis of American Law: Guide to the American Digest System, 2000 Edition, West Group, 1999, which is incorporated herein by reference.) For example, headnote 122 is associated with classes or class identifiers 122.1, 122.2, and 122.3, and headnote 124 is associated with classes or class identifiers 124.1 and 124.2.

In the exemplary system, headnote database 120 includes about 20 million headnotes and grows at an approximate rate of 12,000 headnotes per week. About 89% of the headnotes are associated with a single class identifier, about 10% with two class identifiers, and about 1% with more than two class identifiers.

Additionally, headnote database 120 includes a number of headnotes, such as headnotes 126 and 128, that are not yet assigned or associated with an ALR annotation in database 110. The headnotes, however, are associated with class identifiers. Specifically, headnote 126 is associated with class identifiers 126.1 and 126.2, and headnote 128 is associated with class identifier 128.1.

Coupled to both ALR annotation database 110 and headnote database 120 is classification processor 130. Classification processor 130 includes classifiers 131, 132, 133, and 134, a composite-score generator 135, an assignment decision-maker 136, and decision-criteria module 137. Processor

130 determines whether one or more cases associated with headnotes in headnote database 120 should be assigned to or cited within one or more of the annotations of annotation database 110. Processor 130 is also coupled to preliminary classification database 140.

5 Preliminary classification database 140 stores and/or organizes the assignment or citation recommendations. Within database 140, the recommendations can be organized as a single first-in-first-out (FIFO) queue, as multiple FIFO queues based on single annotations or subsets of annotations. The recommendations are ultimately distributed to work center 150.

10 Work center 150 communicates with preliminary classification database 140 as well as annotation database 110 and ultimately assists users in manually updating the ALR annotations in database 110 based on the recommendations stored in database 140. Specifically, work center 150 includes workstations 152, 154, and 156. Workstation 152, which is substantially identical to workstations 154 and 156, includes a graphical-user interface 152.1, and user-interface devices, such as a keyboard and mouse (not shown.)

15 In general, exemplary system 100 operates as follows. Headnotes database 120 receives a new set of headnotes (such as headnotes 126 and 128) for recently decided cases, and classification processor 130 determines whether one or more of the cases associated with the headnotes are sufficiently relevant to any of the annotations within ALR to justify recommending assignments of the headnotes (or associated cases) to one or more of the annotations. (Some other embodiments directly assign the headnotes or associated cases to the annotations.) The assignment recommendations are stored in preliminary classification database 140 and later retrieved by or presented to editors in work center 150 via graphical-user interfaces in workstations 152, 154, and 156 for acceptance or rejection. Accepted recommendations are added as citations to the respective annotations in ALR annotation database 110 and rejected recommendations are not. However, both accepted and rejected recommendations are fed back to classification processor 130 for incremental training or tuning of its decision criteria.

More particularly, Figure 2 shows a flow chart 200 illustrating in greater detail an exemplary method of operating system 100. Flow chart 200 includes a number of process blocks 210-250. Though arranged serially in the exemplary

embodiment, other embodiments may reorder the blocks, omits one or more blocks, and/or execute two or more blocks in parallel using multiple processors or a single processor organized as two or more virtual machines or subprocessors. Moreover, still other embodiments implement the blocks as one
 5 or more specific interconnected hardware or integrated-circuit modules with related control and data signals communicated between and through the modules. Thus, the exemplary process flow is applicable to software, firmware, hardware, and hybrid implementations.

The remainder of the description uses the following notational system.
 10 The lower case letters a , h , and k respectively denote an annotation, a headnote, and a class or class identifier, such as a West Key Number class or class identifier. The upper case letters A , H , and K respectively denote the set of all annotations, the set of all headnotes, and the set of all key numbers classifications. Additionally, variables denoting vector quantities are in bold-
 15 faced capital letters, and elements of the corresponding vectors are denoted in lower case letters. For example, \mathbf{V} denotes a vector, and v denotes an element of vector \mathbf{V} .

At block 210, the exemplary method begins by representing the annotations in annotations database 110 (in Figure 1) as text-based feature
 20 vectors. In particular, this entails representing each annotation a as a one-column feature vector, \mathbf{V}_a , based on the noun and/or noun-word pairs occurring in headnotes for the cases cited within the annotation. (Other embodiments represent the headnotes as bigrams or noun phrases.)

Although it is possible to use all the headnotes associated with the cases
 25 cited in the annotation, the exemplary embodiment selects from the set of all headnotes associated with the cited cases those that are most relevant to the annotation being represented. For each annotation, this entails building a feature vector using all the headnotes in all cases cited in the annotation and selecting from each case one, two, or three headnotes based on similarity between the
 30 headnotes in a cited case and those of the citing annotation and denoting the most similar headnote(s) as relevant. To determine the most relevant headnotes, the exemplary embodiment uses classifiers 131-134 to compute similarity scores, averages the four scores for each headnote, and defines as most relevant

the highest scoring headnote plus those with a score of at least 80% of the highest score. The 80% value was chosen empirically.

Once selected, the associated headnotes (or alternatively the actual text of the annotations) are represented as a set of nouns, noun-noun, noun-verb, and noun-adjective pairs that it contains. Words in a word-pair are not necessarily adjacent, but are within a specific number of words or characters of each other, that is, within a particular word or character window. The window size is adjustable and can take values from 1 to the total number of words or characters in the headnote. Although larger windows tend to yield better performance, in the exemplary embodiment, no change in performance was observed for windows larger than 32 non-stop words. For convenience, however, the exemplary window size is set to the actual headnote size. The exemplary embodiment excludes stop words and uses the root form of all words. Appendix A shows an exemplary list of exemplary stopwords; however, other embodiments use other lists of stopwords.

Figure 3 shows an example of a headnote 310 and a noun-word representation 320 in accord with the exemplary embodiment. Also shown are West Key Number classification text 330 and class identifier 340.

In a particular annotation vector V_a , the weight, or magnitude, of any particular element v_a is defined as

$$v_a = tf'_a * idf'_a, \quad (1)$$

where tf'_a denotes the term frequency (that is, the total number of occurrences) of the term or noun-word pair associated with annotation a . (In the exemplary embodiment, this is the number of occurrences of the term within the set of headnotes associated with the annotation.) idf'_a denotes the inverse document frequency for the associated term or noun-word pair. idf'_a is defined as

$$idf'_a = \log \left(\frac{N}{df'_a} \right), \quad (2)$$

where N is the total number of headnotes (for example, 20 million) in the collection, and df'_a is the number of headnotes (or more generally documents) containing the term or noun-word pair. The prime ' notation indicates that these frequency parameters are based on proxy text, for example, the text of associated

headnotes, as opposed to text of the annotation itself. (However, other embodiments may use all or portions of text from the annotation alone or in combination with proxy text, such as headnotes or other related documents.)

Even though the exemplary embodiment uses headnotes associated with
5 an annotation as opposed to text of the annotation itself, the annotation-text
vectors can include a large number of elements. Indeed, some annotation
vectors can include hundreds of thousands of terms or noun-word pairs, with the
majority of them having a low term frequency. Thus, not only to reduce the
number of terms to a manageable number, but also to avoid the rare-word
10 problem known to exist in vector-space models, the exemplary embodiment
removes low-weight terms.

Specifically, the exemplary embodiment removes as many low-weight
terms as necessary to achieve a lower absolute bound of 500 terms or a 75%
reduction in the length of each annotation vector. The effect of this process on
15 the number of terms in an annotation vector depends on their weight distribution.
For example, if the terms have similar weights, approximately 75% of the terms
will be removed. However, for annotations with skewed weight distributions, as
few as 10% of the terms might be removed. In the exemplary embodiment, this
process decreased the total number of unique terms for all annotation vectors
20 from approximately 70 million to approximately 8 million terms.

Some other embodiments use other methods to limit vector size. For
example, some embodiments apply a fixed threshold on the number of terms per
category, or on the term's frequency, document frequency, or weight. These
methods are generally efficient when the underlying categories do not vary
25 significantly in the feature space. Still other embodiments perform feature
selection based on measures, such as mutual information. These methods,
however, are computationally expensive. The exemplary method attempts to
strike a balance between these two ends.

Block 220, executed after representation of the annotations as text-based
30 feature vectors, entails modeling one or more input headnotes from database 120
(in Figure 1) as a set of corresponding headnote-text vectors. The input
headnotes include headnotes that have been recently added to headnote database
120 or that have otherwise not previously been reviewed for relevance to the
ALR annotations in database 110.

The exemplary embodiment represents each input headnote h as a vector \mathbf{V}_h , with each element v_h , like the elements of the annotation vectors, associated with a term or noun-word pair in the headnote. v_h is defined as

$$v_h = tf_h * idf_H, \quad (3)$$

- 5 where tf_h denotes the frequency (that is, the total number of occurrences) of the associated term or noun-word pair in the input headnote, and idf_H denotes the inverse document frequency of the associated term or noun-word pair within all the headnotes.

- At block 230, the exemplary method continues with operation of
10 classification processor 130 (in Figure 1). Figure 2 shows that block 230 itself comprises sub-process blocks 231-237.

- Block 231, which represents operation of classifier 131, entails computing a set of similarity scores based on the similarity of text in each input headnote text to the text associated with each annotation. Specifically, the
15 exemplary embodiment measures this similarity as the cosine of the angle between the headnote vector \mathbf{V}_h and each annotation vector \mathbf{V}_a .

Mathematically, this is expressed as

$$S_1 = \cos \theta_{ah} = \frac{\mathbf{V}_a' \cdot \mathbf{V}_h'}{\|\mathbf{V}_a'\| \times \|\mathbf{V}_h'\|}, \quad (4)$$

- where “ \cdot ” denotes the conventional dot- or inner-product operator, and \mathbf{V}_a' and
20 \mathbf{V}_h' denote that respective vectors \mathbf{V}_a and \mathbf{V}_h have been modified to include elements corresponding to terms or noun-word pairs found in both the annotation text and the headnote. In other words, the dot product is computed based on the intersection of the terms or noun-word pairs. $\|X\|$ denotes the length of the vector argument. In this embodiment, the magnitudes are
25 computed based on all the elements of the vector.

- Block 232, which represents operation of classifier 132, entails determining a set of similarity scores based on the similarity of the class identifiers (or other meta-data) associated with the input headnote and those associated with each of the annotations. Before this determination is made, each annotation a is
30 represented as an annotation-class vector \mathbf{V}_a^C vector, with each element

v_a^C indicating the weight of a class identifier assigned to the headnotes cited by the annotation. Each element v_a^C is defined as

$$v_a^C = tf_a^C * idf_a^C, \quad (5)$$

where tf_a^C denotes the frequency of the associated class identifier, and idf_a^C ,
5 denotes its inverse document frequency. idf_a^C is defined as

$$idf_a^C = \log \left(\frac{N_C}{df^C} \right), \quad (6)$$

where N_C is the total number of classes or class identifiers. In the exemplary embodiment, N_C is 91997, the total number of classes in the West Key Number System. df^C is the frequency of the class identifier amongst the set of class
10 identifiers for annotation a . Unlike the exemplary annotation-text vectors which are based on a selected set of annotation headnotes, the annotation-class vectors use all the class identifiers associated with all the headnotes that are associated with the annotation. Some embodiments may use class-identifier pairs, although they were found to be counterproductive in the exemplary implementation.

15 Similarly, each input headnote is also represented as a headnote-class vector V_h^C , with each element indicating the weight of a class or class identifier assigned to the headnote. Each element v_h^C is defined as

$$v_h^C = tf_h^C * idf_h^C, \quad (7)$$

with tf_h^C denoting the frequency of the class identifier, and idf_h^C denoting the
20 inverse document frequency of the class identifier. idf_h^C is defined as

$$idf_h^C = \log \left(\frac{N_C}{df_a^C} \right), \quad (8)$$

where N_C is the total number of classes or class identifiers and df_h is the frequency of the class or class identifier amongst the set of class or class identifiers associated with the annotation.

25 Once the annotation-class and headnote-class vectors are established, classification processor 130 computes each similarity score S_2 as the cosine of the angle between them. This is expressed as

$$S_2 = \cos \theta_{ah} = \frac{\mathbf{V}_a^C \cdot \mathbf{V}_h^C}{\|\mathbf{V}_a^C\| \times \|\mathbf{V}_h^C\|}, \quad (9)$$

For headnotes that have more than one associated class identifier, the exemplary embodiment considers each class identifier separately of the others for that headnote, ultimately using the one yielding the maximum class-identifier

5 similarity. The maximization criteria is used because, in some instances, a headnote may have two or more associated class identifiers (or Key Number classifications), indicating its discussion of two or more legal points. However, in most cases, only one of the class identifiers is relevant to a given annotation.

In block 233, classifier 133 determines a set of similarity scores S_3 based
10 on the probability that a headnote is associated with a given annotation from class-identifier (or other meta-data) statistics. This probability is approximated by

$$S_3 = P(h | a) = P(\{k\}_h | a) = \max_{k' \in \{k\}_h} P(k' | a), \quad (10)$$

where $\{k\}_h$ denotes the set of class identifiers assigned to headnote h . Each

15 annotation conditional class probability $P(k | a)$ is estimated by

$$P(k | a) = \frac{1 + tf_{(k,a)}}{|a| + \sum_{k' \in a} tf_{(k',a)}}, \quad (11)$$

where $tf_{(k,a)}$ is the term frequency of the k -th class identifier among the class identifiers associated with the headnotes of annotation a ; $|a|$ denotes the total

number of unique class identifiers associated with annotation a (that is, the
20 number of samples or cardinality of the set); and $\sum_{k' \in a} tf_{(k',a)}$ denotes the sum of the term frequencies for all the class identifiers.

The exemplary determination of similarity scores S_3 relies on assumptions that class identifiers are assigned to a headnote independently of each other, and that only one class identifier in $\{k\}_h$ is actually relevant to
25 annotation a . Although the one-class assumption does not hold for many annotations, it improves the overall performance of the system.

Alternatively, one can multiply the conditional class-identifier (Key Number classifications) probabilities for the annotation, but this effectively

penalizes headnotes with multiple Key Number classifications (class assignments), compared to those with single Key Number classifications. Some other embodiments use Bayes' rule to incorporate *a priori* probabilities into classifier 133. However, some experimentation with this approach suggests that system performance is likely to be inferior to that provided in this exemplary implementation.

The inferiority may stem from the fact that annotations are created at different times, and the fact that one annotation has more citations than another does not necessarily mean it is more probable to occur for a given headnote. Indeed, a greater number of citations might only reflect that one annotation has been in existence longer and/or updated more often than another. Thus, other embodiments might use the prior probabilities based on the frequency that class numbers are assigned to the annotations.

In block 234, classifier 134 determines a set of similarity scores S_4 , based on $P(a | h)$, the probability of each annotation given the text of the input headnote. In deriving a practical expression for computing $P(a | h)$, the exemplary embodiment first assumes that an input headnote h is completely represented by a set of descriptors T , with each descriptor t assigned to a headnote with some probability, $P(t | h)$. Then, based on the theory of total probability and Bayes' theorem, $P(a | h)$ is expressed as

$$\begin{aligned} P(a | h) &= \sum_{t \in T} P(a | h, t) P(t | h) \\ &= \sum_{t \in T} \frac{P(h | a, t) P(a | t)}{P(h | t)} P(t | h) \end{aligned} \quad (12)$$

Assuming that a descriptor is independent of the class identifiers associated with a headnote allows one to make the approximation:

$$P(h | a, t) \approx P(h | t) \quad (13)$$

and to compute the similarity scores S_4 according to

$$S_4 = P(a | h) = \sum_{t \in T} P(t | h) P(a | t) \quad (14)$$

where $P(t | h)$ is approximated by

$$P(t | h) = \frac{tf_{(t,h)}}{\sum_{t' \in T} tf_{(t',h)}}. \quad (15)$$

$tf_{(t,h)}$ denotes the frequency of term t in the headnote and $\sum_{t' \in T} tf_{(t',h)}$ denotes the sum of the frequencies of all terms in the headnote. $P(a | t)$ is defined according to Bayes' theorem as

$$P(a | t) = \frac{P(t | a) P(a)}{\sum_{a' \in A} P(t | a') P(a')}, \quad (16)$$

where $P(a)$ denotes the prior probability for annotation a , and $P(t | a)$, the probability of a discriminator t given annotation a , is estimated as

$$P(t | a) \cong \frac{1}{|a|} \sum_{h \in a} P(t | h), \quad (17)$$

and $\sum_{a' \in A}$ denotes summation over all annotations a' in the set of annotations A .

- 10 Since all the annotation prior probabilities $P(a)$ and $P(a')$ are assumed to be equal, $P(a | t)$ is computed using

$$P(a | t) = \frac{P(t | a)}{\sum_{a' \in A} P(t | a')}. \quad (18)$$

- Block 235, which represents operation of composite-score generator 135, entails computing a set of composite similarity scores CS_a^h based on the sets of similarity scores determined at blocks 231-235 by classifiers 131-135, with each composite score indicating the similarity of the input headnote h to each annotation a . More particularly, generator 135 computes each composite score CS_a^h according to

$$CS_a^h = \sum_{i=1}^4 w_{ia} S_{a,i}^h, \quad (19)$$

- 20 where $S_{a,i}^h$ denotes the similarity score of the i -th similarity score generator for the input headnote h and annotation a , and w_{ia} is a weight assigned to the i -th similarity score generator and annotation a . Execution of the exemplary method then continues at block 236.

At block 236, assignment decision-maker 136 recommends that the input headnote or a document, such as a case, associated with the headnote be classified or incorporated into one or more of the annotations based on the set of composite scores and decision criteria within decision-criteria module 137. In the exemplary embodiments, the headnote is assigned to annotations according to the following decision rule:

$$\text{If } CS_a^h > \Gamma_a, \text{ then recommend assignment of } h \text{ or } D_h \text{ to annotation } a, \quad (20)$$

where Γ_a is an annotation-specific threshold from decision-criteria module 137 and D_h denotes a document, such as a legal opinion, associated with the headnote. (In the exemplary embodiment, each ALR annotation includes the text of associated headnotes and its full case citation.)

The annotation-classifier weights w_{ia} , for $i = 1$ to 4, $a \in A$, and the annotation thresholds Γ_a , $a \in A$, are learned during a tuning phase. The weights, $0 \leq w_{ia} \leq 1$, reflect system confidence in the ability of each similarity score to route to annotation a . Similarly, the annotation thresholds Γ_a , $a \in A$, are also learned and reflect the homogeneity of an annotation. In general, annotations dealing with narrow topics tend to have higher thresholds than those dealing with multiple related topics.

In this ALR embodiment, the thresholds reflect that, over 90% of the headnotes (or associated documents) are not assigned to any annotations. Specifically, the exemplary embodiment estimates optimal annotation-classifier weights and annotation thresholds through exhaustive search over a five-dimensional space. The space is discretized to make the search manageable. The optimal weights are those corresponding to maximum precision at recall levels of at least 90%.

More precisely, this entails trying every combination of four weight variables, and for each combination, trying 20 possible threshold values over the interval $[0,1]$. The combination of weights and threshold that yields the best precision and recall is then selected. The exemplary embodiment excludes any weight-threshold combinations resulting in less than 90% recall.

To achieve higher precision levels, the exemplary embodiment effectively requires assignments to compete for their assigned annotations or target classifications. This competition entails use of the following rule:

$$\text{Assign } h \text{ to } a, \text{ iff } CS_a^h > \alpha \hat{S} \quad (21)$$

- 5 where α denotes an empirically determined value greater than zero and less than 1, for example, 0.8; \hat{S} denotes the maximum composite similarity score associated with a headnote in $\{H_a\}$, the set of headnotes assigned to annotation a .

Block 240 entails processing classification recommendations from
10 classification processor 130. To this end, processor 130 transfers classification recommendations to preliminary classification database 140 (shown in Figure 1). Database 140 sorts the recommendation based on annotation, jurisdiction, or other relevant criteria and stores them in, for example, a single first-in-first-out (FIFO) queue, as multiple FIFO queue based on single annotations or subsets of
15 annotations.

One or more of the recommendations are then communicated by request or automatically to workcenter 150, specifically workstations 152, 154, and 156. Each of the workstations displays, automatically or in response to user activation, one or more graphical-user interfaces, such as graphical-user interface
20 152.1.

Figure 4 shows an exemplary form of graphical-user interface 152.1. Interface 152.1 includes concurrently displayed windows or regions 410, 420, 430 and buttons 440-490.

Window 410 displays a recommendation list 412 of headnote identifiers
25 from preliminary classification database 140. Each headnote identifier is logically associated with at least one annotation identifier (shown in window 430). Each of the listed headnote identifiers is selectable using a selection device, such as a keyboard or mouse or microphone. A headnote identifier 412.1 in list 412 is automatically highlighted, by for example, reverse-video
30 presentation, upon selection. In response, window 420 displays a headnote 422 and a case citation 424, both of which are associated with each other and the highlighted headnote identifier 412.1. In further response, window 430 displays at least a portion or section of an annotation outline 432 (or classification

hierarchy), associated with the annotation designated by the annotation identifier associated with headnote 412.1.

Button 440, labeled “New Section,” allows a user to create a new section or subsection in the annotation outline. This feature is useful, since in some instances, a headnote suggestion is good, but does not fit an existing section of the annotation. Creating the new section or subsection thus allows for convenient expansion of the annotation..

Button 450 toggles on and off the display of a text box describing headnote assignments made to the current annotation during the current session. In the exemplary embodiment, the text box presents each assignment in a short textual form, such as <annotation or class identifier><subsection or section identifier ><headnote identifier>. This feature is particularly convenient for larger annotation outlines that exceed the size of window 430 and require scrolling contents of the window.

Button 460, labeled “Un-Allocate,” allows a user to de-assign, or declassify, a headnote to a particular annotation. Thus, if a user changes her mind regarding a previous, unsaved, classification, the user can nullify the classification. In some embodiments, headnotes identified in window 410 are understood to be assigned to the particular annotation section displayed in window 430 unless the user decides that the assignment is incorrect or inappropriate. (In some embodiments, acceptance of a recommendation entails automatic creation of hyperlinks linking the annotation to the case and the case to the annotation.)

Button 470, labeled “Next Annotation,” allows a user to cause display of the set of headnotes recommended for assignment to the next annotation. Specifically, this entails not only retrieving headnotes from preliminary classification storage 140 and displaying them in window 410, but also displaying the relevant annotation outline within window 430.

Button 480, labeled “Skip Anno,” allows a user to skip the current annotation and its suggestions altogether and advance to the next set of suggestions and associated annotation. This feature is particularly useful when an editor wants another editor to review assignments to a particular annotation, or if the editor wants to review this annotation at another time, for example, after reading or studying the entire annotation text, for example. The suggestions

remain in preliminary classification database 140 until they are either reviewed or removed. (In some embodiments, the suggestions are time-stamped and may be supplanted with more current suggestions or deleted automatically after a preset period of time, with the time period, in some variations dependent on the particular annotation.)

Button 490, labeled "Exit," allows an editor to terminate an editorial session. Upon termination, acceptances and recommendations are stored in ALR annotations database 110.

Figure 2 shows that after processing of the preliminary classifications, execution of the exemplary method continues at block 250. Block 250 entails updating of classification decision criteria. In the exemplary embodiment, this entails counting the numbers of accepted and rejected classification recommendations for each annotation, and adjusting the annotation-specific decision thresholds and/or classifier weights appropriately. For example, if 80% of the classification recommendations for a given annotation are rejected during one day, week, month, quarter or year, the exemplary embodiment may increase the decision threshold associated with that annotation to reduce the number of recommendations. Conversely, if 80% are accepted, the threshold may be lowered to ensure that a sufficient number of recommendations are being considered.

Exemplary System for Classifying Headnotes to American Jurisprudence

Figure 5 shows a variation of system 100 in the form of an exemplary classification system 500 tailored to facilitate classification of documents to one or more of the 135,000 sections of The American Jurisprudence (AmJur). Similar to an ALR annotation, each AmJur section cites relevant cases as they are decided by the courts. Likewise, updating AmJur is time consuming.

In comparison to system 100, classification system 500 includes six classifiers: classifiers 131-134 and classifiers 510 and 520, a composite score generator 530, and assignment decision-maker 540. Classifiers 131-134 are identical to the ones used in system 100, with the exception that they operate on AmJur data as opposed to ALR data.

Classifiers 510 and 520 process AmJur section text itself, instead of proxy text based on headnotes cited within the AmJur section.. More

specifically, classifier 510 operates using the formulae underlying classifier 131 to generate similarity measurements based on the tf-idfs (term-frequency-inverse document frequency) of noun-word pairs in AmJur section text. And classifier 520 operates using the formulae underlying classifier 134 to generate similarity measurements based on the probabilities of a section text given the input headnote.

Once the measurements are computed, each classifier assigns each AmJur section a similarity score based on a numerical ranking of its respective set of similarity measurements. Thus, for any input headnote, each of the six classifiers effectively ranks the 135,000 AmJur sections according to their similarities to the headnote. Given the differences in the classifiers and the data underlying their scores, it is unlikely that all six classifiers would rank the most relevant AmJur section the highest; differences in the classifiers and the data they use generally suggest that this will not occur. Table 1 shows a partial ranked listing of AmJur sections showing how each classifier scored, or ranked, their similarity to a given headnote.

Section	C 1 Ranks	C 2 Ranks	C 3 Ranks	C 4 Ranks	C 5 Ranks	C 6 Ranks	Median Ranks
Section_1	1	8	4	1	3	2	2.5
Section_2	3	2	5	9	1	3	3
Section_3	2	4	6	5	4	4	4
Section_4	5	1	3	8	6	1	4
Section_5	7	3	2	2	5	5	4
Section_6	4	5	1	7	2	9	4.5
Section_7	8	7	8	4	7	6	7
Section_8	6	9	7	3	10	7	7
Section_9	9	10	9	6	9	10	9
Section_10	10	6	10	10	8	8	9

Table 1: Partial Ranked Listing AmJur Sections
based of Median of Six Similarity Scores

Composite score generator 530 generates a composite similarity score for each AmJur section based on its corresponding set of six similarity scores. In the exemplary embodiment, this entails computing the median of the six scores for each AmJur section. However, other embodiments can compute a uniform or non-uniformly weighted average of all six or a subset of the six rankings. Still other embodiments can select the maximum, minimum, or mode as the composite score for the AmJur section. After generating the composite scores, the composite score generator forwards data identifying the AmJur section associated with the highest composite score, the highest composite score, and the input headnote to assignment decision-maker 540.

Assignment decision-maker 540 provides a fixed portion of headnote-classification recommendations to preliminary classification database 140, based on the total number of input headnotes per a fixed time period. The fixed number and time period governing the number of recommendations are determined according to parameters within decision-criteria module 137. For example, one embodiment ranks all incoming headnotes for the time period, based on their composite scores and recommends only those headnotes that rank in the top 16 percent.

In some instances, more than one headnote may have a composite score that equals a given cut-off threshold, such as top 16%. To ensure greater accuracy in these circumstances, the exemplary embodiment re-orders all headnote-section pairs that coincide with the cut-off threshold, using the six actual classifier scores.

This entails converting the six classifier scores for a particular headnote-section pair into six Z-scores and then multiplying the six Z-scores for a particular headnote-section pair to produce a single similarity measure. (Z-scores are obtained by assuming that each classifier score has a normal distribution, estimating the mean and standard deviation of the distribution, and then subtracting the mean from the classifier score and dividing the result by the standard deviation.) The headnote-section pairs that meet the acceptance criteria are then re-ordered, or re-ranked, according to this new similarity measure, with as many as needed to achieve the desired number of total recommendations being forwarded to preliminary classification database 140. (Other embodiments may apply this "reordering" to all of the headnote-section pairs and then filter

these based on the acceptance criteria necessary to obtain the desired number of recommendations.)

Exemplary System for Classifying Headnotes to West Key Number System

- 5 Figure 6 shows another variation of system 100 in the form of an exemplary classification system 600 tailored to facilitate classification of input headnotes to classes of the West Key Number System. The Key Number System is a hierarchical classification system with 450 top-level classes, which are further subdivided into 92,000 sub-classes, each having a unique class identifier.
- 10 In comparison to system 100, system 600 includes classifiers 131 and 134, a composite score generator 610, and an assignment decision-maker 620.

- In accord with previous embodiments, classifiers 131 and 134 model each input headnote as a feature vector of noun-word pairs and each class identifier as a feature vector of noun-word pairs extracted from headnotes
- 15 assigned to it. Classifier 131 generates similarity scores based on the tf-idf products for noun-word pairs in headnotes assigned to each class identifier and to a given input headnote. And classifier 134 generates similarity scores based on the probabilities of a class identifier given the input headnote. Thus, system 600 generates over 184,000 similarity scores, with each scores representing the
- 20 similarity of the input headnote to a respective one of the over 92,000 class identifiers in the West Key Number System using a respective one of the two classifiers.

- Composite score generator 610 combines the two similarity measures for each possible headnote-class-identifier pair to generate a respective composite
- 25 similarity score. In the exemplary embodiment, this entails defining, for each class or class identifier, two normalized cumulative histograms (one for each classifier) based on the headnotes already assigned to the class. These histograms approximate corresponding cumulative density functions, allowing one to determine the probability that a given percentage of the class identifiers
- 30 scored below a certain similarity score.

 More particularly, the two cumulative normalized histograms for class-identifier c , based on classifiers 131 and 134 are respectively denoted F_c^1 and F_c^2 , and estimated according to:

$$F_c^1(s) = F_c^1(s - 0.01) + \frac{1}{M_c} * \left| \left\{ h_i \mid S_i^1 = s \right\} \right| \quad (22)$$

and

$$F_c^2(s) = F_c^2(s - 0.01) + \frac{1}{M_c} * \left| \left\{ h_i \mid S_i^2 = s \right\} \right|, \quad (23)$$

where c denotes a particular class or class identifier;

- 5 $s = 0, 0.01, 0.02, 0.03, \dots, 1.0$; $F(s < 0) = 0$; M_c denotes the number of headnotes classified to or associated with class or class identifier c ; $|B|$ denotes the number of elements in the set B ; $h_i, i = 1, \dots, M_c$ denotes the set of headnotes already classified or associated with class or class identifier c ; S_i^1 denotes the similarity score for headnote h_i and class-identifier c , as measured by classifier 131, and S_i^2 denote the similarity score for headnote h_i and class-identifier c , as measured by classifier 134. (In this context, each similarity score indicates the similarity of a given assigned headnote to all the headnotes assigned to class c .) In other words, $\left| \left\{ h_i \mid S_i^1 = s \right\} \right|$ denotes the number of headnotes assigned to class c that received a score of s from classifier 131, and $\left| \left\{ h_i \mid S_i^2 = s \right\} \right|$ denotes the number of headnotes assigned to class c that received a score of s from classifier 134.

- Thus, for every possible score value (between 0 and 1 with a particular score spacing), each histogram provides the percentage of assigned headnotes that scored higher and lower than that particular score. For example, for classifier 131, the histogram for class identifier c might show that 60% of the set of headnotes assigned to classifier c scored higher than 0.7 when compared to the set of headnotes as a whole; whereas for classifier 134 the histogram might show that 50% of the assigned headnotes scored higher than 0.7

- Next, composite score generator 610 converts each score for the input headnote into a normalized similarity score using the corresponding histogram and computes each composite score for each class based on the normalized scores. In the exemplary embodiment, this conversion entails mapping each classifier score to the corresponding histogram to determine its cumulative probability and then multiplying the cumulative probabilities of respective pairs

of scores associated with a given class c to compute the respective composite similarity score. The set of composite scores for the input headnote are then processed by assignment decisionmaker 620.

- Assignment decision maker 620 forwards a fixed number of the top
- 5 scoring class identifiers to preliminary classification database 140. The exemplary embodiments suggest the class identifiers having the top five composite similarity scores for every input headnote.

Other Exemplary Applications

The components of the various exemplary systems presented can be combined in myriad ways to form other classification systems of both greater and lesser complexity. Additionally, the components and systems can be tailored for other types of documents other than headnotes. Indeed, the components and systems and embodied teachings and principles of operation are relevant to virtually any text or data classification context.

For example, one can apply one or more of the exemplary systems and related variations to classify electronic voice and mail messages. Some mail classifying systems may include one or more classifiers in combination with conventional rules which classify messages as useful or SPAM based on whether the sender is in your address book, same domain as recipient, etc.

Appendix AExemplary Stop Words

a a.m ab about above accordingly across ad after afterward afterwards again
 against ago ah ahead ain't all allows almost alone along already alright also
 although always am among amongst an and and/or anew another ante any
 anybody anybody's anyhow anymore anyone anyone's anything anything's
 anytime anytime's anyway anyways anywhere anywhere's anywise appear
 approx are aren't around as aside associated at available away awfully awhile b
 banc be became because become becomes becoming been before beforehand
 behalf behind being below beside besides best better between beyond both brief
 but by bythe c came can can't cannot cant cause causes certain certainly cetera cf
 ch change changes cit cl clearly cmt co concerning consequently consider
 contain containing contains contra corresponding could couldn't course curiam
 currently d day days dba de des described di did didn't different divers do does
 doesn't doing don't done down downward downwards dr du during e e.g each ed
 eds eg eight eighteen eighty either eleven else elsewhere enough especially et etc
 even ever evermore every everybody everybody's everyone everyone's
 everyplace everything everything's everywhere everywhere's example except f
 facie facto far few fewer fide fides followed following follows for forma former
 formerly forth forthwith fortiori fro from further furthermore g get gets getting
 given gives go goes going gone got gotten h had hadn't happens hardly has hasn't
 have haven't having he he'd he'll he's hello hence henceforth her here here's
 hereabout hereabouts hereafter herebefore hereby herein hereinafter hereinbefore
 hereinbelow hereof hereto heretofore hereunder hereunto hereupon herewith hers
 herself hey hi him himself his hither hitherto hoc hon how howbeit however
 howsoever hundred i i'd i'll i'm i've i.e ibid ibidem id ie if ignored ii iii illus
 immediate in inasmuch inc indeed indicate indicated indicates infra initio insofar
 instead inthe into intra inward ipsa is isn't it it's its itself iv ix j jr judicata just k
 keep kept kinda know known knows l la last later latter latterly le least les less
 lest let let's like likewise little looks ltd m ma'am many may maybe me
 meantime meanwhile mero might million more moreover most mostly motu mr
 mrs ms much must my myself name namely naught near necessary neither never
 nevermore nevertheless new next no no-one nobody nohow nolo nom non none

1003794-12301
 1003794-12301

nonetheless noone nor normally nos not nothing novo now nowhere o o'clock of
 ofa off ofhis oft often ofthe ofthis oh on once one one's ones oneself only onthe
 onto op or other others otherwise ought our ours ourself ourselves out outside
 over overall overly own p p.m p.s par para paras pars particular particularly
 passim per peradventure percent perchance perforce perhaps pg pgs placed
 please plus possible pp probably provides q quite r rata rather really rel relatively
 rem res resp respectively right s sa said same says se sec seem seemed seeming
 seems seen sent serious several shall shalt she she'll she's should shouldn't since
 sir so some somebody somebody's somehow someone someone's something
 something's sometime sometimes somewhat somewhere somewhere's specified
 specify specifying still such sundry sup t take taken tam than that that's thats the
 their theirs them themselves then thence thenceforth thenceforward there there's
 thereafter thereby therefor therefore therefrom therein thereof thereon theres
 thereto theretofore thereunto thereupon therewith these they they'll thing things
 third this thither thorough thoroughly those though three through throughout thru
 thus to to-wit together too toward towards u uh unless until up upon upward
 upwards used useful using usually v v.s value various very vi via vii viii
 virtually vs w was wasn't way we we'd we'll we're we've well went were weren't
 what what'll what's whatever whatsoever when whence whenever where
 whereafter whereas whereat whereby wherefore wherefrom wherein whereinto
 whereof whereon wheresoever whereto whereunder whereunto whereupon
 wherever wherewith whether which whichever while whither who who'd who'll
 who's whoever whole wholly wholly whom whose why will with within without
 won't would wouldn't x y y'all ya'll ye yeah yes yet you you'll you're you've your
 yours yourself yourselves z

10027944 " 4324104
 10027944 " 4324104

Conclusion

In furtherance of the art, the inventors have presented various exemplary systems, methods, and software which facilitate the classification of text, such as headnotes or associated legal cases to a classification system, such as that represented by the nearly 14,000 ALR annotations. The exemplary system classifies or makes classification recommendations based on text and class similarities and probabilistic relations. The system also provides a graphical-user interface to facilitate editorial processing of recommended classifications and thus automated update of document collections, such as the American Legal Reports, American Jurisprudence, and countless others.

The embodiments described above are intended only to illustrate and teach one or more ways of practicing or implementing the present invention, not to restrict its breadth or scope. The actual scope of the invention, which embraces all ways of practicing or implementing the teachings of the invention, is defined only by the following claims and their equivalents.

1003794-139101